

Microsoft 70-454

**Upgrade: Transition Your MCITP SQL Server 2005 DBD
to MCITP SQL Server 2008 DBD**

Version: 34.0

QUESTION NO: 1

You have a table that has 10 million rows. The table has a view that returns all of the rows.

You discover performance degradation when you run an unfiltered view.

You need to recommend a solution to replace the view. The solution must require that the returned data is filtered by a parameter.

What should you use?

- A. an indexed view
- B. a scalar function
- C. a table-valued function
- D. a table-valued type

Answer: C

Explanation:

QUESTION NO: 2

You are a database developer. You plan to design a database solution by using SQL Server 2008.

There are two schemas named Sales and Marketing. You are the owner of the Sales schema and the Marketing schema is owned by a user named MarketingUser.

Users of the Marketing schema do not have permissions to access the Sales schema. You have permissions to create objects in all schemas in the database.

The Sales schema has a table named Customers.

You plan to create a stored procedure in the Marketing schema for the marketing team. The stored procedure will select data from the Customers table and will be owned by MarketingUser.

You need to ensure that the marketing team is able to execute the stored procedure.

What should you do?

- A. Create the procedure by using the EXECUTE AS SELF option.
- B. Create the procedure by using the EXECUTE AS CALLER option.
- C. Create the procedure by using the EXECUTE AS OWNER option.
- D. Create the procedure by using the EXECUTE AS USER=MarketingUser option.

Answer: A

Explanation:

QUESTION NO: 3

You have a table named Table1. A sample of the data in Table1 is shown in the following table.

SalesID	SalesOrderNumber
1	SO2159
2	SO2768
3	SO3978
4	SO3010
5	SO4818
6	SO3919
7	SO3999

There is a defined nonclustered index on the SalesOrderNumber column. The following query executes against the table.

```
WHERE 'S03' = LEFT (SalesOrderNumber,3)
```

You need to minimize the amount page I/O that is generated when the query runs.

What should you do?

- A. Use a query hint.
- B. Add a non-filtered index.
- C. Rewrite the WHERE clause to use a LIKE statement.
- D. Rewrite the WHERE clause to use a substring function on the SalesOrderNumber column.

Answer: C

Explanation:

QUESTION NO: 4

You need to design a stored procedure that contains a transaction. If an error condition causes the transaction to roll back, a temporary result set must be available to the query that follows the failed transaction.

Which temporary storage object should you use?

- A. a common table expression (CTE)
- B. a derived table
- C. a table variable
- D. a temporary table

Answer: C

Explanation:

QUESTION NO: 5

You are a database developer. You plan to design a database solution by using SQL Server 2008.

A frequently used query takes very long to execute.

You discover that the query frequently uses full-table scans instead of indexes. This causes other queries that modify the table to be blocked.

The indexing strategy on the underlying tables that the query uses can change.

You need to design a solution that performs the following tasks:

- Removes full-table scans
- Allows the query optimizer to select the appropriate index.

What should you do?

- A. Use the INDEX table hint.
- B. Use the INDEX(O) table hint.
- C. Use the NOEXPAND table hint.
- D. Use the FORCESEEK table hint.

Answer: D

Explanation:

QUESTION NO: 6

You are a database developer. You create a database by using SQL Server 2008 in an enterprise environment.

The database contains two stored procedures named ModifySales and RetrieveSalesTrend. The

ModifySales stored procedure uses a transaction that updates a table named SalesOrders. The RetrieveSalesTrend stored procedure retrieves and aggregates data from the SalesOrders table for a sales trend analysis report.

Both stored procedures are executed frequently each day.

Users report a considerable wait time while they run the sales trend analysis report.

You need to ensure that sales trend analysis report runs as quickly as possible.

What should you do?

- A. Change the isolation level to SERIALIZABLE for ModifySales.
- B. Change the isolation level to READ UNCOMMITTED for ModifySales.
- C. Add the NOWAIT hint to the SELECT statement in RetrieveSalesTrend.
- D. Add the NOLOCK hint to the SELECT statements in RetrieveSalesTrend.

Answer: D

Explanation:

QUESTION NO: 7

You are a database developer. You plan to design a database solution by using SQL Server 2008.

The database application has a table named Transactions that contains millions of rows. The table has multiple columns that include transaction_id and transaction_date. There is a clustered index on the transaction_id column. There is a nonclustered index on the transaction_date column. You discover that the following query takes a long time to execute.

```
SELECT transaction_id, transaction_date, transaction_notes
FROM transactions
WHERE transaction_type_id = 'FXO'
AND transaction_date between @start_date and @end_date
```

The summary of the execution plan is as shown in the following code segment.

```
--Filter(WHERE: {[transaction_type_id]='FXO'})
--Nested Loops(Inner Join)
--Index Seek(OBJECT: {[transactions]. [nc_transactions_transaction_date])
--Clustered Index Seek(OBJECT: {[transactions]. [PK_transactions_transaction_id])
```

You need to ensure that the query retrieves data in minimum possible time.

What should you do?

- A. Create a nonclustered index on the transaction_type_id column.
- B. Create a nonclustered index on the transaction_date and transaction_type_id columns.
- C. Create a nonclustered index on the transaction_date column and include the transaction_type_id and transaction_notes columns.
- D. Create a nonclustered index on the transaction_date and transaction_type_id columns and include the transaction_notes column.

Answer: D

Explanation:

QUESTION NO: 8

You have a table named Books that contains information about books. Books has the columns in the following table.

Column	Data type	Primary key	Indexed	Index type
BookId	int	Yes	Yes	nonclustered
Title	nvarchar(50)	No	No	-
ISBN	nvarchar(10)	No	No	-
Description	nvarchar(400)	No	No	-

You plan to create several queries that will filter on Title and ISBN. The queries will return values from Title, ISBN, and Description. You need to recommend an indexing solution to meet the following requirements:

- Minimize the amount of time required to return the results of the planned queries
- Minimize the number of indexes

What should you recommend?

- A. Create a nonclustered index on each column.
- B. Create a clustered index on Title, ISBN and Description as the key value.
- C. Create a clustered index on Title and ISBN and set the index fill factor to 75.
- D. Create a nonclustered index on Title and ISBN and include the Description column.

Answer: D

Explanation:

QUESTION NO: 9

You are designing a table that will contain the columns in the following table,

Column	Description
Invoice#	Unique number
Subtotal	Subtotal of sale
TaxAmt	Taxes
FreightAmt	Shipping charges

You need to add a column that will display the sum of the Subtotal, TaxArnt, and FreightAmt columns. The solution must minimize the amount of storage space required to store the table.

What should you create?

- A. a column that uses the money data type
- B. a column that uses the numeric data type
- C. a computed column
- D. a persisted computed column

Answer: C

Explanation:

QUESTION NO: 10

You have a database that is used for archiving and reporting.

The database is currently 100 GB and increases in size at the rate of 10 percent per month. A SQL Server Profiler trace shows that high levels of disk I/O create a performance bottleneck. You need to optimize the database to meet the following requirements:

- Reduce the disk I/O required by database
- Reduce the storage requirements for the database

What should you do?

- A. Enable row compression.
- B. Implement table partitioning.
- C. Run the DBCC SHRINKDATABASE command.
- D. Run the DBCC SHRINKFILE command.

Answer: A

Explanation:

QUESTION NO: 11

You need to design a table that will contain a monetary value. The value must support five digits to the right of the decimal point.

Which data type should you choose?

- A. bigint
- B. decimal
- C. int
- D. money

Answer: B

Explanation:

QUESTION NO: 12

You need to develop a stored procedure strategy that will manipulate temporary result sets to meet the following requirements:

- Stored procedures must receive result sets as parameters
- The strategy must be developed by using the minimum amount of effort

The result sets will contain less than 1,000 rows.

What should you use?

- A. a common table expression (CTE)
- B. a scalar function
- C. a user-defined table type
- D. a view

Answer: C

Explanation:

QUESTION NO: 13

You are a database developer. You provide solutions by using SQL Server 2008 in an enterprise environment.

Your online transaction processing (OLTP) database contains a table named SalesOrders, Your data warehouse contains a table named factBuyingHabits. The factBuyingHabits table has no indexes.

You need to synchronize data between the two tables on a weekly basis. The synchronization process has the following requirements:

- New records in the SalesOrders table are inserted in the factBuyingHabits table.
- When a record is modified in the SalesOrders table, the modification is updated in the factBuyingHabits table.
- Records that are deleted from the SalesOrders table are also deleted from the factBuyingHabits table.

You need to design an appropriate synchronization solution. You want to achieve this goal by using minimum amount of coding and administrative efforts.

What should you do?

- A.** Design an SSIS package each for the INSERT, UPDATE, and DELETE operations. Schedule a job to run this package.
- B.** Design a single SSIS package that uses the Slowly changing Dimension task, schedule a job to run this package.
- C.** Write one stored procedure that contains a MERGE statement to perform the INSERT, UPDATE, and DELETE operations. Schedule a job to run the stored procedure.
- D.** Write three stored procedures each for the INSERT, UPDATE, and DELETE operations. Schedule a job to run the stored procedures in a sequential manner.

Answer: C

Explanation:

QUESTION NO: 14

You are a database developer. You plan to design a database solution by using SQL Server 2008,

The database has a table named Sales. The Sales table contains 10 million rows.

You discover that the following query takes a long time to execute.

```
SELECT s.sale_id, ...
FROM Sales AS s
JOIN Country AS c
ON s.Country_id = c.Country_id
AND c.Country_name = 'USA'
```

A summary of the execution plan is as shown in the following code segment.

```
--Hash Match(Inner Join, HASH: ([s].[Country_id]) = ([c].[Country_id])
--Clustered Index Scan(OBJECT: ([Country].[PK_Country_Country_id] AS [c])
--Clustered Index Scan(OBJECT: ([Sales].[PK_Sales_Sale_id] AS [s]))
```

You need to ensure that the query retrieves data in minimum possible time.

What should you do?

- A. Modify the query to use a loop join hint.
- B. Modify the query to use a merge join hint.
- C. Create a nonclustered index in the Country_id column of the Sales table.
- D. Create a nonclustered index in the Country_name column of the Country table.

Answer: C

Explanation:

QUESTION NO: 15

You are a database developer. You plan to design a database solution by using SQL Server 2008.

A database contains a table that has a column defined as a smallint data type. The table is partitioned on the basis of the smallint column. T partition function, named MyRangePFI, is a range-right function that has boundaries of 100 and 1,000.

The table must be altered to contain the following partitions: